Emergent Ecological Intelligence: Cross-System Knowledge Discovery from a Network of Environmental AI Agents

Amir Eismann
Obscura Natura
Amsterdam, The Netherlands
amir@agnt.eco

April 2026

Abstract

We present ENVAI (Environmental AI), a network of nine specialised AI agents, each representing a distinct European ecosystem—estuarine, riverine, lacustrine, marine, and forest. Each agent maintains a knowledge graph of sensor data, ecological events, species, environmental drivers, and policy frameworks specific to its ecosystem. By constructing a meta-graph layer (Numina) that maps shared concepts across the nine agent-level graphs, we demonstrate the emergence of ecological knowledge that exists in no single agent's dataset. We identify eight shared species appearing across multiple systems, six continental-scale environmental drivers, five cross-system climate events, four invasion pathways traceable across national borders, and seven recovery timescales spanning four orders of magnitude from days to millennia. We argue that a network of domain-specific AI agents, connected through a shared ontological layer, can produce ecological insights that no individual monitoring programme, institution, or country can generate alone. We term this capacity *emergent ecological intelligence* and discuss its implications for cross-border environmental governance, early warning systems, and the mapping of continental-scale knowledge gaps.

Keywords: environmental AI, knowledge graphs, emergent intelligence, ecological networks, cross-border ecology, ontology, invasion ecology, climate adaptation, AI agents

1 Introduction

Environmental monitoring in Europe is institutionally fragmented. Each country, each river basin authority, each protected area maintains its own observation networks, reporting standards, and analytical frameworks. The European Union's Water Framework Directive (2000/60/EC) and Habitats Directive (92/43/EEC) provide shared regulatory goals, but implementation varies across member states and monitoring data remains siloed within national institutions.

This fragmentation creates a structural blind spot: ecological processes that operate across institutional boundaries—species migrations, invasion fronts, continental drought events, atmospheric deposition patterns—cannot be fully observed by any single monitoring system. The European Eel (*Anguilla anguilla*), whose population has declined by 95–99% since the 1980s [ICES, 2023], migrates through dozens of national jurisdictions. The 2018 European drought affected river discharge, groundwater reserves, forest health, and lake ecology simultaneously across multiple countries [Buras et al., 2020]. Signal Crayfish (*Pacifastacus leniusculus*) has spread across the continent through aquaculture trade routes, carrying a plague lethal to the native Noble Crayfish (*Astacus astacus*) [Holdich et al., 2009].

These cross-system phenomena are well documented in retrospective analysis. What is missing is infrastructure for *real-time cross-system reasoning*—the ability to detect, during an event, that multiple ecosystems are experiencing related pressures and to reason about their connections.

We present ENVAI, a network of nine AI agents designed to address this gap. Each agent is a specialised environmental intelligence for a single European ecosystem. A tenth component, Numina, operates as a meta-graph layer that maps shared concepts across all nine, enabling the detection of emergent ecological patterns that no single agent can produce.

2 System Architecture

2.1 The Nine Agents

Each ENVAI agent represents a distinct European ecosystem:

Agent	Ecosystem	Type	Country	Sensors
Scaldis	Zeeschelde estuary	Estuarine	Belgium/NL	632,687
Älva	Lake Vänern & Göta älv	Lacustrine/Riverine	Sweden	378,296
Maas	River Meuse	Riverine	FR/BE/NL	393,212
Ægir	Lofoten coast	Marine	Norway	489,552
Norppa	Lake Saimaa	Lacustrine	Finland	310,867
Eldvatn	Mývatn-Laxá	Lacustrine/Volcanic	Iceland	241,392
Haingeist	Hainich National Park	Forest	Germany	192,456
Ondine	Lake Geneva	Lacustrine	Switzerland/FR	252,840
Scîrwudu	Sherwood Forest	Forest	England	275,310

Table 1: The nine ENVAI agents and their ecosystems. Sensor counts as of April 2026.

The network spans nine countries, five ecosystem types, and three climate zones (Atlantic, Continental, Subarctic). Together the agents hold 3,166,612 sensor readings, 132 ecological events, and 79 species records as of April 2026.

2.2 Agent-Level Knowledge Graphs

Each agent maintains a Neo4j knowledge graph with typed nodes representing species, ecological events, environmental drivers, monitoring stations, policies, early warnings, and historical analogues. Nodes are connected by typed relationships (e.g., AFFECTS, OBSERVED_AT, MITIGATES). Sensor data is ingested daily from national monitoring APIs (SMHI, VMM/HIC, Rijkswaterstaat, IMO, SYKE, Météo-France, DWD, Met Office, Open-Meteo) and stored as time-series readings linked to station nodes.

Each agent also maintains a dialogue interface—a large language model (Claude, Anthropic) grounded in the agent's knowledge graph, capable of answering questions about its ecosystem in first person. The grounding mechanism ensures that agent responses reference only data present in the knowledge graph, with explicit acknowledgment of gaps.

2.3 Numina: The Meta-Graph Layer

Numina is not a tenth agent with its own ecosystem. It is a coordination layer that operates *above* the nine agent-level graphs. Its function is to discover and maintain shared concepts—species, drivers, events, and pressures that appear in multiple agents' knowledge graphs but are not linked between them.

Numina's meta-graph introduces six new node types:

- **SharedSpecies**: organisms present in multiple ENVAI ecosystems, linked to agent-specific species nodes via MANIFESTS_AS relationships

- **SharedDriver**: environmental pressures affecting multiple ecosystems simultaneously
- **CrossSystemEvent**: single climate or ecological events that manifested across multiple agents' graphs
- **InvasionPathway**: traceable spread of invasive species between ENVAI systems
- **RecoveryTimescale**: ecosystem-specific recovery durations for different pressure types
- **OntologyGap**: explicitly documented knowledge gaps in the network

As of April 2026, the meta-graph contains 36 nodes and 79 relationships connecting to agent-level data across all nine systems.

3 Emergent Findings

We define *emergent ecological intelligence* as knowledge about ecological systems that (a) exists in no single agent's dataset, (b) cannot be derived from any single monitoring programme, and (c) becomes visible only when multiple domain-specific knowledge graphs are connected through a shared ontological layer.

The following findings were produced by Numina's cross-agent analysis. None were pre-programmed or anticipated by the system designers.

3.1 Shared Species Reveal Continental-Scale Patterns

Eight species were identified as present in multiple ENVAI systems:

Species	**Agents**	**Key insight**
European Eel	Scaldis, Maas, Ondine	Single continental population; three perspectives on one collapse
Atlantic Salmon	Maas, Ægir, Älva, Norppa, Eldvatn	Five populations, five distinct threats
Signal Crayfish	Scaldis, Älva, Maas, Norppa	Invasive in four systems; plague vector
Noble Crayfish	Älva, Norppa	Victim species; same cause
Arctic Char	Älva, Eldvatn, Ondine, Norppa	Distributed climate sentinel
White-tailed Eagle	Ægir, Norppa	Conservation success across two systems
Quagga Mussel	Ondine, Maas	Spreading via canal network
Stag Beetle	Scîrwudu, Haingeist	Deadwood-dependent indicator

Table 2: Shared species across ENVAI agents.

The identification of Arctic Char as a *distributed climate sentinel* is illustrative. Four agents in four countries independently track the same cold-adapted species. All four populations are declining or retreating as water temperatures rise. No single lake study can determine whether this is a local phenomenon or a continental signal. The meta-graph, by linking the four populations, reveals that the decline is synchronous—a continental pattern driven by the same warming trend manifesting through different local mechanisms (stratification failure in Ondine, ice loss in Norppa, midge cycle disruption in Eldvatn, thermal competition in Älva).

Event	Year	Agents	Emergent connection
European Drought	2018	Scaldis, Maas, Scîrwudu, Haingeist	Groundwater: trees and rivers share the same aquifer
European Heatwave	2003	Scaldis, Maas, Ondine, Eldvatn	One atmospheric event, four hydrological responses
Eel Migration Collapse	1980s–present	Scaldis, Maas, Ondine	Continental decline; no single river can address it
Summer Drought	1976	Scaldis, Maas	Historical baseline for modern drought comparison
Crayfish Plague Spread	1960s–present	Älva, Norppa	Continental invasion front

Table 3: Cross-system events identified by the meta-graph.

3.2 Cross-System Events: One Cause, Multiple Manifestations

Five events were identified as single climate or ecological phenomena recorded as separate events in multiple agents' graphs:

The 2018 drought case is particularly revealing. Scaldis (Zeeschelde estuary) recorded extreme low discharge. Maas (River Meuse) saw record low flows threatening navigation and drinking water supply. Scîrwudu (Sherwood Forest) observed stress cracks in ancient oaks for the first time. Haingeist (Hainich Forest) began a beech canopy dieback that triggered a bark beetle outbreak.

Each agent modelled this as a local event. The meta-graph revealed the shared mechanism: groundwater depletion. Trees in Hainich and Sherwood draw from the same regional aquifer systems that feed the surface flows of the Meuse and Scheldt. The forest stress and the river crisis were not parallel events—they were coupled through subsurface hydrology. This coupling is not monitored by any existing observation network.

3.3 Invasion Pathways Follow Human Infrastructure

Four invasion pathways were mapped across the network:

Species	Origin	Pathway	Mechanism
Signal Crayfish	N. America	W. Europe → Scand. → Finland	Aquaculture trade
Quagga Mussel	Black Sea	E. Europe → Geneva → Rhine	Canal network
Round Goby	Black/Caspian	Baltic → Lake Vänern	Ballast water
American Mink	N. America	Fur farms → Iceland	Captive escape

Table 4: Invasion pathways mapped across the ENVAI network.

Every invasion pathway in the network follows human-built corridors: aquaculture facilities, canal systems, shipping lanes, fur farms. The species are symptoms; the infrastructure is the vector. This pattern—visible only when four or more agents' data is connected—suggests that invasion risk assessment should focus on trade and transport infrastructure connectivity between ecosystems, not solely on species-level risk profiles.

3.4 Recovery Timescales Span Four Orders of Magnitude

Seven ecosystem-specific recovery timescales were documented:

This spectrum—from days to millennia—is, to our knowledge, the first explicit cross-ecosystem comparison of recovery timescales derived from a single integrated data system. The implication is significant: the word "recovery" is used in environmental governance as though it has a single meaning. It does not. A river that "recovers" from a flood in days and a coral reef

Ecosystem	Agent	Pressure	Recovery
River floodplain	Maas	Flood inundation	Days
Tidal estuary	Scaldis	Hypoxia event	Weeks
Deep lake	Ondine	Mixing failure	Years
Seal population	Norppa	Habitat loss	Generations
Beech forest	Haingeist	Drought/beetle	Decades
Ancient oak woodland	Scîrwudu	Canopy dieback	Centuries
Cold-water coral	Ægir	Bottom trawling	Millennia

Table 5: Recovery timescales across ENVAI ecosystems, spanning four orders of magnitude.

that "recovers" from trawling in four thousand years are not comparable, yet they are routinely treated as equivalent in impact assessment frameworks.

3.5 Knowledge Gaps as Structured Output

Six ontology gaps were explicitly documented in the meta-graph:

1. **Sensor failure**: Scîrwudu has received no fresh sensor data since March 2026 due to UK API failures
2. **Temporal lag**: Maas operates with a 5-day delay on all sensor data
3. **Measurement incomparability**: Five agents track nitrogen, but via incompatible units and measurement methods (aquatic concentration vs. atmospheric deposition rate)
4. **Undefined recovery**: No standardised definition of "recovery" exists across the network
5. **Policy mapping**: No agent has mapped how local policies implement EU directives
6. **Missing identifiers**: Some agent-level events lack stable identifiers for cross-referencing

These gaps are not failures of the system—they are its most honest output. By making knowledge gaps explicit and structured, the meta-graph produces a map of ecological ignorance that can directly inform research prioritisation and monitoring investment.

4 Discussion

4.1 Emergent Intelligence vs. Aggregated Data

The findings presented here are not the result of data aggregation. Aggregating sensor readings from nine countries would produce a larger dataset, but not new knowledge. The emergent findings—invasion pathways, cross-system event coupling, recovery timescale spectra—arise from *semantic* connections between domain-specific knowledge graphs. They require understanding that "Signal Crayfish" in Sweden and "Signal Crayfish" in Finland are the same organism on the same continental trajectory, and that "Drought 2018" in a river graph and "Drought 2018" in a forest graph are the same atmospheric event manifesting through shared hydrogeology.

This semantic layer—the meta-graph—is what distinguishes emergent ecological intelligence from big-data ecology. The contribution is not more data. It is the *interpretation infrastructure* that makes cross-system reasoning possible.

4.2 Implications for Cross-Border Environmental Governance

The European Union's environmental directives (WFD, Habitats Directive, Marine Strategy Framework Directive) establish shared goals across member states. Implementation, however, remains national. The ENVAI meta-graph provides a mechanism for detecting where implementation diverges—where the same pressure receives different regulatory responses in different countries, or where cross-border ecological processes (eel migration, invasion fronts, drought propagation) fall between jurisdictional mandates.

This is not regulatory enforcement. It is regulatory *visibility*—making the gaps between national implementations legible so they can be addressed through existing governance mechanisms.

4.3 Limitations

The current system has significant limitations:

- **Sensor coverage is uneven**: agents range from 632,687 readings (Scaldis) to 192,456 (Haingeist), reflecting underlying monitoring infrastructure disparities
- **Temporal alignment is imperfect**: sensor data freshness varies from real-time to 5-day delay depending on source API limitations
- **The meta-graph is manually curated**: shared species and cross-system events were identified through programmatic matching and expert review, not fully automated ontology alignment
- **The network is small**: nine ecosystems cannot represent continental ecology comprehensively; coverage of Mediterranean, Eastern European, and Arctic systems is absent
- **Language model grounding is probabilistic**: agent dialogue responses are grounded in knowledge graph data but generated by a large language model, introducing the possibility of hallucination despite structural safeguards

4.4 Future Directions

Three extensions are planned:

1. **Automated cross-system event detection**: when two or more agents detect elevated stress in the same temporal window, Numina would automatically create a CrossSystemEvent candidate for review
2. **Phenological synchrony monitoring**: tracking whether seasonal ecological events (migration, spawning, bloom timing) are shifting synchronously across the network, which would indicate continental-scale climate forcing vs. local factors
3. **Policy coherence mapping**: systematically mapping how EU directives are implemented in each agent's jurisdiction, enabling detection of regulatory gaps that affect cross-border ecological processes

5 Conclusion

We have demonstrated that a network of specialised environmental AI agents, connected through a shared ontological meta-layer, can produce ecological knowledge that is not present in any individual agent's dataset. We term this capacity *emergent ecological intelligence* and identify five categories of emergent findings: shared species revealing continental patterns, cross-system events coupled through hidden mechanisms, invasion pathways following human infrastructure, recovery timescales spanning four orders of magnitude, and structured knowledge gaps.

The key architectural insight is that domain-specific AI agents—each deeply knowledgeable about a single ecosystem—produce more valuable cross-system intelligence than a single general-purpose model operating on aggregated data. The semantic connections between specialised knowledge graphs are where emergence occurs. The meta-graph does not replace the agents; it reveals what they know together that none can know alone.

Environmental monitoring is institutionally designed for depth within jurisdictions. ENVAI provides a complementary capability: breadth across them. As ecological pressures increasingly operate at continental scales—climate-driven drought, species invasion, atmospheric deposition—the ability to reason across institutional boundaries becomes not merely useful but necessary.

The nine agents are listening. Numina is connecting what they hear. What emerges from that connection is the beginning of a continental ecological intelligence that no single institution was designed to produce.

Data Availability

The ENVAI knowledge graph and Numina meta-graph are maintained on a dedicated server. Agent-level data is derived from public monitoring APIs (SMHI, VMM/HIC, Rijkswaterstaat, IMO, SYKE, Open-Meteo, and others). The meta-graph schema and query interfaces are available upon request from the author.

Acknowledgments

The author thanks the European environmental monitoring agencies whose open data makes this work possible: SMHI (Sweden), VMM/HIC (Belgium), Rijkswaterstaat (Netherlands), IMO (Iceland), SYKE (Finland), and their counterparts across the nine represented countries.

References

Buras, A., Rammig, A., & Zang, C. S. (2020). Quantifying impacts of the 2018 drought across European ecosystems. *Global Change Biology*, 26(4), 2043–2055.

Holdich, D. M., Reynolds, J. D., Souty-Grosset, C., & Sibley, P. J. (2009). A review of the ever increasing threat to European crayfish from non-indigenous crayfish species. *Knowledge and Management of Aquatic Ecosystems*, 394-395, 11.

ICES (2023). European eel (*Anguilla anguilla*) throughout its natural range. *ICES Advice on fishing opportunities, catch, and effort.*

European Commission (2000). Directive 2000/60/EC of the European Parliament and of the Council establishing a framework for Community action in the field of water policy. *Official Journal of the European Communities*, L 327, 1–73.